

Removal of Batch Effects using Distribution-Matching Residual Networks

Uri Shaham^{1*}, Kelly P. Stanton^{2*}, Jun Zhao³, Huamin Li⁴, Khadir Raddassi⁵,
Ruth Montgomery⁶, and Yuval Kluger^{2,3,4†}

December 26, 2016

¹Department of Statistics, Yale University, New Haven, CT, USA

²Department of Pathology, Yale School of Medicine, New Haven, CT, USA

³Program of Computational Biology and Bioinformatics, Yale University, New Haven, CT, USA

⁴Applied Mathematics Program, Yale University, New Haven, CT, USA

⁵Departments of Neurology and Immunobiology, Yale School of Medicine, New Haven, CT, USA

⁶Department of Internal Medicine, Yale School of Medicine, New Haven, CT, USA

Abstract

Sources of variability in experimentally derived data include measurement error in addition to the physical phenomena of interest. This measurement error is a combination of systematic components, originating from the measuring instrument, and random measurement errors. Several novel biological technologies, such as mass cytometry and single-cell RNA-seq, are plagued with systematic errors that may severely affect statistical analysis if the data is not properly calibrated. We propose a novel deep learning approach for removing systematic batch effects. Our method is based on a residual network, trained to minimize the Maximum Mean Discrepancy (MMD) between the multivariate distributions of two replicates, measured in different batches. We apply our method to mass cytometry and single-cell RNA-seq datasets, and demonstrate that it effectively attenuates batch effects.

1 Introduction

Biological data are affected by the conditions of the measuring instruments. For example, biomedical data from replicated¹ measurements, measured in different batches, may be distributed differently due to variation in these conditions between batches. The term *batch effects*, often used in the biological community, describes a situation where subsets (batches) of

*The first two authors contributed equally to this work.

†To whom correspondence should be addressed, yuval.kluger@yale.edu

¹We use the term *replicates* to refer to technical replicates, i.e., multiple measurements of the same specimen, for example, two blood drops of the same person.

the measurements significantly differ in distribution, due to irrelevant instrument-related factors (Leek et al., 2010). Batch effects introduce systematic error, which may cause statistical analysis to produce spurious results and/or obfuscate the signal of interest.

For example, CyTOF, a mass cytometry technique for measuring multiple protein levels in many cells of a biological specimen, is known to incur batch effects. When replicate blood specimens from the same patient are measured on a CyTOF machine in different batches (e.g. different days), they might differ noticeably in the distribution of cells in the multivariate protein space. In order to run a valid and effective statistical analysis on the data, a calibration process has to be carried out, to account for the effect of the difference in instrument conditions on the measurements.

Typically, the systematic effect of varying instrument conditions on the measurements depends on many unknown factors, whose impact on the difference between the observed and underlying true signal cannot be modeled. In this manuscript, we consider cases where replicates differ in distribution, due to batch effects. By designating one replicate to be the source sample² and the other to be the target sample, we propose a deep learning approach to learn a map that calibrates the distribution of the source sample to match that of the target. Our proposed approach is designed for data where the difference between these source and target distributions is moderate, so that the map that calibrates them is close to the identity map; such an assumption is fairly realistic in many situations. An example of the problem and the output of our proposed method is depicted in Figure 1. A short demo movie is available at <https://www.youtube.com/watch?v=Lqya9WDkZ60>.

To evaluate the effectiveness of our proposed approach, we employ it to analyze CyTOF and single-cell RNA-seq, and demonstrate that it successfully removes significant portions of the batch effect. We also demonstrate that it outperforms other popular approaches for calibration. To the best of our knowledge, similar performance on CyTOF data has never been reported.

We justify our approach by demonstrating that shortcut connections are crucial to the success of calibration, as nets without shortcut connections might distort the biological properties of the data. Furthermore, we provide evidence that a map from a source replicate to a target replicate, both extracted from the same specimen, can also be learned in an indirect manner, through maps between replicates from other specimens, without losing much accuracy. Generalizing this approach allows one to calibrate multiple source batches to a single target batch, where replicates from a single reference specimen are measured in each of these batches.

The remainder of this manuscript is organized as follows: in Section 2 we give a brief review of Maximum Mean Discrepancy and Residual Nets, on which our approach is based. The calibration learning problem is defined in Section 3, where we also describe our proposed approach. Experimental results on CyTOF and single-cell RNA-seq measurements are reported in Section 4. In Section 5 we review some related works. In Section 6 we discuss some technical aspects of our approach. Section 7 concludes the manuscript.

²The term *sample* is used with different meanings in the biological and statistical communities. Both meanings are used in this manuscript, however, usage should be clear from context. .




Figure 1: Calibration of CyTOF data. Projection of the source (red) and target (blue) samples on the first two principal components of the target data. Left: before calibration. Right: after calibration.

2 Preliminaries

2.1 Maximum Mean Discrepancy

Maximum Mean Discrepancy (MMD, Gretton et al. (2012, 2006)) is a measure for distance between two probability distributions p, q . It is defined with respect to a function class \mathcal{F} by

$$\text{MMD}(\mathcal{F}, p, q) \equiv \sup_{f \in \mathcal{F}} (\mathbb{E}_{x \sim p} f(x) - \mathbb{E}_{x \sim q} f(x)). \quad (1)$$

When \mathcal{F} is a reproducing kernel Hilbert space with kernel k , the MMD can be written as the distance between the mean embeddings of p and q

$$\text{MMD}^2(\mathcal{F}, p, q) = \|\mu_p - \mu_q\|_{\mathcal{F}}^2, \quad (2)$$

where $\mu_p(t) = \mathbb{E}_{x \sim p} k(x, t)$. Equation (2) can be written as

$$\text{MMD}^2(\mathcal{F}, p, q) = \mathbb{E}_{x, x' \sim p} k(x, x') - 2\mathbb{E}_{x \sim p, y \sim q} k(x, y) + \mathbb{E}_{y, y' \sim q} k(y, y'), \quad (3)$$

where x and x' are independent, and so are y and y' . Importantly, if k is a universal kernel, then $\text{MMD}(\mathcal{F}, p, q) = 0$ iff $p = q$. In practice, the distributions p, q are unknown, and instead we are given observations $X = \{x_1, \dots, x_n\}, Y = \{y_1, \dots, y_m\}$, so that the (biased) sample version

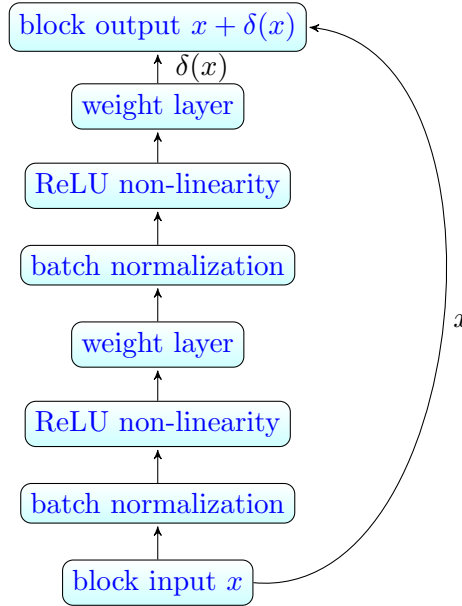


Figure 2: A typical ResNet block.

of (3) becomes

$$\text{MMD}^2(\mathcal{F}, X, Y) = \frac{1}{n^2} \sum_{x_i, x_j \in X} k(x_i, x_j) - \frac{2}{nm} \sum_{x_i \in X, y_j \in Y} k(x_i, y_j) + \frac{1}{m^2} \sum_{y_i, y_j \in Y} k(y_i, y_j).$$

MMD was originally proposed as a non-parametric two sample test, and has since been widely used in various applications. Li et al. (2015); Dziugaite et al. (2015), use it as a loss function for neural net; here we adopt this direction to tackle the calibration problem, as discussed in Section 3.

2.2 Residual Nets

Residual neural networks (ResNets), proposed by He et al. (2015) and improved in (He et al., 2016), is a recently introduced class of very deep neural nets, mostly used for image recognition tasks. ResNets are typically formed by concatenation of many blocks, where each block receives an input x (the output of the previous block) and computes output $y = x + \delta(x)$, where $\delta(x)$ is the output of a small neural net, which usually consists of two sequences of batch normalization (Ioffe and Szegedy, 2015), weight layers and non-linearity activations, as depicted in Figure 2.

It was empirically shown by He et al. (2015) that the performance of very deep convolutional nets without shortcut connections deteriorates beyond some depth, while ResNets can

grow very deep with increasing performance. In a subsequent work, He et al. (2016) showed that the gradient backpropagation in ResNets is improved, by avoiding exploding or vanishing gradients, comparing to networks without shortcut connections; this allows for more successful optimization, regardless of the depth. Li et al. (2016) showed that ResNets with shortcut connections of depth 2 are easy to train, while deeper shortcut connections make the loss surface more flat. In addition, they argue that initializing ResNets with weights close to zero performs better than other standard initialization techniques.

Since a ResNet block consists of a residual term and an identity term, it can easily learn functions close to the identity function, when the weights are initialized close to zero, which is shown to be a valuable property for deep neural nets (Hardt and Ma, 2016). In our case, the ability to efficiently learn functions which are close to the identity is appealing from an additional reason: we are interested in performing calibration between replicate samples whose multivariate distributions are close to each other; to calibrate the samples, we are therefore interested in learning a map which is close to the identity map. A ResNet structure is hence a convenient tool to learn such a map.

3 Tackling the Calibration Problem

Formally, we consider the following learning problem: let $\mathcal{D}_1, \mathcal{D}_2$ be two distributions on \mathbb{R}^d , such that there exists a continuous map $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ so that if $X \sim \mathcal{D}_1$ then $\psi(X) \sim \mathcal{D}_2$. We also assume that ψ is a small perturbation of the identity map.

We are given two finite samples $\{x_1, \dots, x_n\}, \{y_1, \dots, y_m\}$ from $\mathcal{D}_1, \mathcal{D}_2$, respectively. The goal is to learn a map $\hat{\psi} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ so that $\{\hat{\psi}(x_1), \dots, \hat{\psi}(x_n)\}$ is likely to be a sample from \mathcal{D}_2 .

Since we assume that ψ is close to the identity, it is convenient to express it as $\psi(x) = x + \delta(x)$, where $\delta(x)$ is small, so that the connection to ResNets blocks becomes apparent.

Our proposed solution, which we term MMD-ResNet is therefore a ResNet; the network gets two samples $\{x_1, \dots, x_n\}, \{y_1, \dots, y_m\}$ of points in \mathbb{R}^d . We refer to $\{x_1, \dots, x_n\}$ as the source sample and to $\{y_1, \dots, y_m\}$ as the target sample. The net receives $\{x_1, \dots, x_n\}$ as input and is trained to learn a map of the source sample, to make it similar in distribution to the target sample. Specifically, we train the net with the following loss function

$$L(w) = \sqrt{\text{MMD}^2(\{\hat{\psi}(x_1), \dots, \hat{\psi}(x_n)\}, \{y_1, \dots, y_m\})},$$

where $\hat{\psi}$ is the map computed by the network, and depends on the network parameters w . We train the net in a stochastic mode, so that in fact the MMD is computed only on mini-batches from both samples, and not on the entire samples.

4 Experimental Results

In this section we report experimental results on biological data obtained using two types of high-throughput technologies: CyTOF and single-cell RNA-seq (scRNA-seq). CyTOF is a mass

cytometry technology that allows simultaneous measurements of multiple protein markers in each cell of a specimen (e.g., a blood sample), consisting of $10^4 - 10^6$ cells (Spitzer and Nolan, 2016). scRNA-seq is a sequencing technology that allows to simultaneously measure mRNA expression levels of all genes in thousands of single cells.

4.1 Technical Details

All MMD-ResNets were trained using RMSprop (Tieleman and Hinton, 2012), using the Keras default hyper-parameter setting; a penalty of 0.01 on the l_2 norm of the network weights was added to the loss for regularization. We used mini-batches of size 1000 from both the source and target samples. A subset 10% of the training data was held out for validation, to determine when to stop the training.

The kernel we used is a sum of three Gaussian kernels

$$k(x, y) = \sum_i \exp \left(\frac{\|x - y\|^2}{\sigma_i^2} \right).$$

We chose the σ_i s to be $\frac{m}{2}, m, 2m$, where m is the median of the average distance between a point in the target sample to its nearest 25 neighbors.

We implemented our net in Keras; our codes and data are publicly available at <https://github.com/ushaham/BatchEffectRemoval.git>.

4.2 Calibration of CyTOF Data

Mass cytometry uses a set of antibodies, each of which is conjugated to a unique heavy ion and binds to a different cellular protein. Cells are then individually nebulized and subjected to mass spectrometry. Protein abundance is indirectly observed from the signal intensity at each protein’s associated ions’ mass to charge ratio. Multiple specimens can be run in the same batch by using barcoding with additional ions to record the origin of each specimen (Spitzer and Nolan, 2016). A CyTOF batch contains measurements of numerous cells from a few specimens, and each batch is affected by systematic errors (Finck et al., 2013).

4.2.1 Data

Our calibration experiments were performed on data collected at Yale New Haven Hospital; Peripheral Blood Mononuclear Cells (PBMC) were collected from two MS patients at baseline and 90 days after Gilenya treatment and cryopreserved. At the end of the study PBMC were thawed in two batches (on two different days) and incubated with or without PMA+ionomycin (using a robotic platform). PMA/ionomycin stimulated and unstimulated samples were barcoded using Cell-ID (Fluidigm), then pooled and labeled for different markers with mass cytometry antibodies and analyzed on CyTOF III Helios. Altogether we used a collection containing eight samples: 2 patients \times 2 conditions \times 2 days. From this collection, we assembled four source-target pairs, where for each patient and biological condition, the sample from day 1 was

considered as source and the one from day 2 as target. All samples were of dimension $d = 25$ ³ and contained 1800-5000 cells.

4.2.2 Pre-processing

All samples were manually filtered by a human expert to remove debris and dead cells. Log transformation, a standard practice for CyTOF, was applied to the data. In addition, a bead-normalization procedure was applied to the data; this is a current practice for normalizing CyTOF data (Finck et al., 2013). Yet, our results demonstrate that the samples clearly differ in distribution, despite the fact that they were normalized.

A typical CyTOF sample contains large proportions of zero values (up to 40% sometimes) which occur due to instabilities of the CyTOF instrument and usually do not reflect biological phenomenon. As leaving the zero values in place might incur difficulties to calibrate the data, a cleaning procedure has to be carried out. In our experiments we collected the cells with non or very few zero values and used them to train a denoising autoencoder (DAE Vincent et al. (2008)). Specifically, the DAE was trained to reconstruct clean cells x from noisy inputs \tilde{x} , where the \tilde{x} was obtained from x by multiplying each entry of x by a independent Bernoulli random variable with parameter = 0.8. The DAE contained two hidden layers, each of 25 ReLU units; the output units were linear. As with the MMD-ResNets, the DAEs were also trained using RMSprop, and their loss contained l_2 penalization of the weights. Once the DAE was trained, we passed the source and target samples through it, and used their reconstructions, which did not contain zeros, for the calibration. In all our CyTOF experiments, source and target refer to the denoised version of these samples. Lastly, as a standard practice, in each of the experiment the input to the net (i.e., the source sample) was standardized to have zero mean and unit variance in each dimension. The parameters of the standardization were then also applied to the target sample.

4.2.3 CyTOF calibration

We trained a MMD-ResNet on each of the four source-target pairs. All nets were identical, and contained three blocks, where each block is as in Figure 2. Each of the weight matrices was of size 25×25 . The net weights were initialized by sampling from a $\mathcal{N}(0, 10^{-4})$ distribution. The projection of the target and source data onto the first two PCs of the target sample in a representative source-target pair is shown in Figure 1. The plots of the remaining three pairs are presented in Appendix B. In the left plot, it is apparent that before calibration, the source sample (red) differs in distribution from the target sample (blue). After calibration (right plot), the gap between the source and the target distributions decreases significantly. The MMD between the source and target before and after calibration in each of the four pairs is shown in Table 1, in addition to the MMD obtained using a multi-layer perceptron (MLP) MMD-net with a similar architecture to the ResNet, except without shortcut connections. The MLP was initialized in a standard fashion (Glorot and Bengio, 2010). As can be seen, the

³See a full specification of the markers in Appendix A

Table 1: CyTOF calibration experiment: MMD values between random batches of size 1000 from the source and target samples, before and after calibration on each of the four source-target pairs (patient1-baseline, patient2-baseline, patient1-treatment, patient2-treatment). The MMD between two random batches of the target sample is provided as reference in the bottom row. The calibrated data is significantly closer in MMD to the target sample. The presented values are average \pm std, based on sampling of five random subsets of size 1000.

MMD to target \ pair	pa.1 base.	pa.2 base.	pa.1 treat.	pa.2 treat.
no calibration	0.66 \pm 0.01	0.56 \pm 0.01	0.59 \pm 0.01	0.70 \pm 0.01
MLP calibration	0.55 \pm 0.01	0.18 \pm 0.01	0.26 \pm 0.01	0.21 \pm 0.01
ResNet calibration	0.27 \pm 0.01	0.17 \pm 0.01	0.24 \pm 0.01	0.17 \pm 0.01
MMD(target,target)	0.12 \pm 0.01	0.12 \pm 0.01	0.13 \pm 0.01	0.13 \pm 0.01

calibrated data is significantly closer to the target data than the original source data. The ResNet achieves similar performance to the MLP on two pairs and outperforms the MLP on the other two. In Section 4.2.4 we will show that ResNet architecture is in fact a crucial element in our approach, for a more important reason.

On a per-marker level, Figure 3 shows the empirical cumulative distribution functions of the first six markers in the source sample before and after the calibration, in comparison to the target sample. In all cases, as well as on the remaining markers that are not shown here, the calibrated source curves are substantially closer to the target than the curves before calibration.

4.2.4 Biological Validation and the Importance of Shortcut Connections

To biologically assess the quality of the calibration and further justify our proposed approach, we inspect the effect of calibration not only at a global level across all types of cell sub-populations, but also zoom in to a specific cell sub-population. Specifically, we focus here on CD8+T-cells, also known as Killer T-cells, in the 2D space of the markers CD28 and GzB. In each sample, we identified the CD8+T-cells sub-population based on manual gating, performed by a human expert. Figure 4 shows the CD8+T-cells of the source and target samples from the baseline samples of patient 2 (patient2-baseline), before calibration, after calibration using ResNet and after calibration using similar net without shortcut connections (MLP). As can be seen, when the calibration is performed by a net without shortcut connections, the CD8+T-cells sub-population is not mapped to the same region as its target sample counterpart. However, with ResNet it is mapped appropriately.

The MMD score between the target sample and the ResNet-calibrated source sample was very similar to the MMD score between the target sample and the MLP-calibrated source sample. We therefore see that in order to achieve good calibration it does not suffice that the calibrated source sample will be close in MMD to the target sample. It is also crucial that the calibration map will be close to the identity. Nets without shortcut connections can clearly

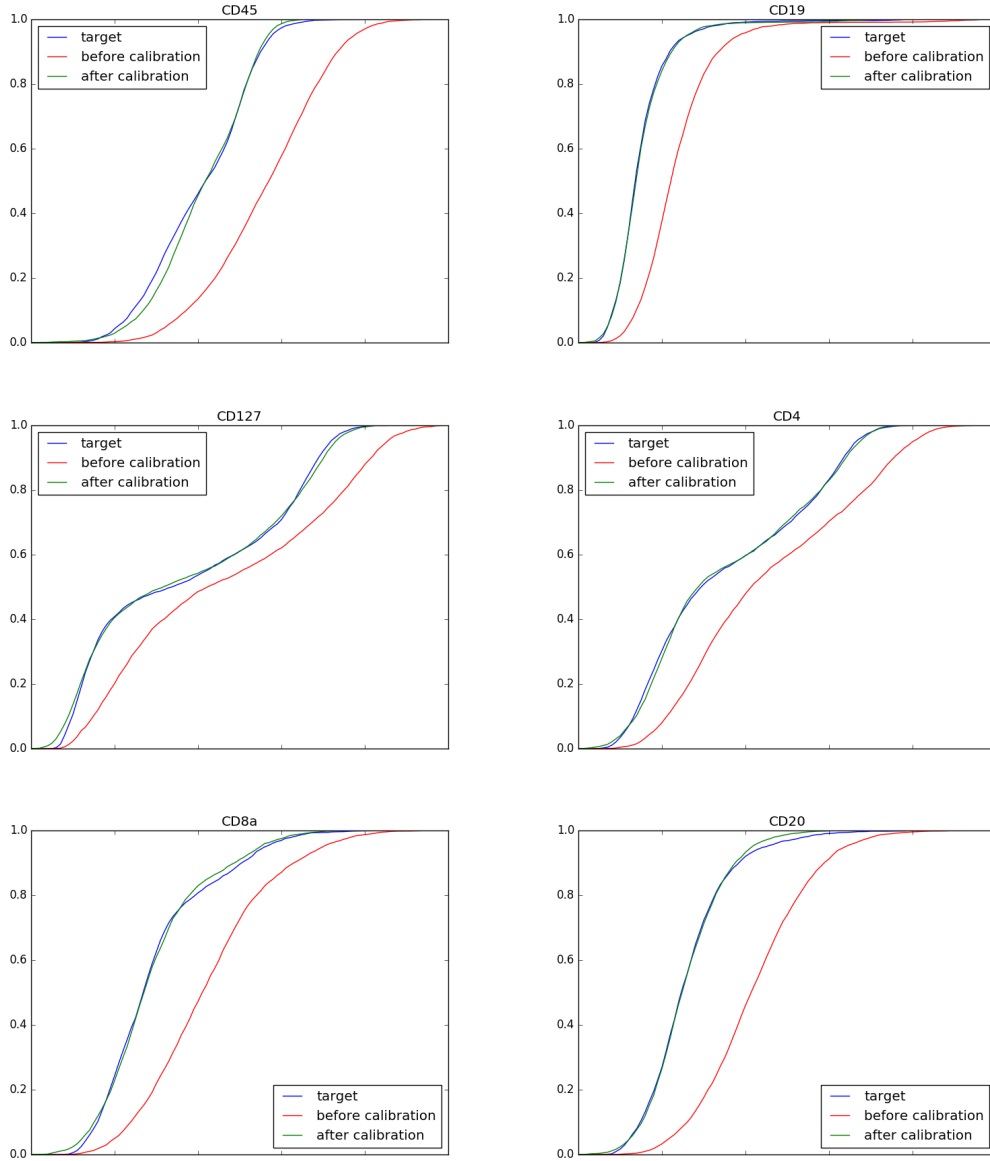


Figure 3: A marginal perspective on the quality of calibration. Empirical cumulative distribution functions of the first six markers in the CyTOF calibration experiment. In each plot the blue, red and green curves corresponds to the target, source and calibrated source samples, respectively. In each marker the blue and green curves are substantially closer than the blue and red curves.




Figure 4: Calibration of CyTOF data: CD8+T-cells (red) and target (blue) samples in the (CD28, GzB) plane. Left: before calibration. Center: calibration using MLP. Right: calibration using ResNet.

compute maps which are close to the identity. However, when trained to minimize MMD, the resulting map is not necessarily close to the identity, as there might be different maps that yield low MMD, despite being far from the identity, and are easier to reach from random initialization by optimization. Therefore, to obtain a map that is close to the identity, ResNet is a more appropriate tool, if not crucial, comparing to nets without shortcut connections.

The plots for the remaining three source-target pairs are shown in Appendix B.

4.2.5 Comparison to Linear Methods

In this section we compare the quality of calibration of our MMD-ResNet to two of the most popular techniques for removal of batch effects. The simplest and most common (Nygaard et al., 2016) adjustment is zero centering, i.e., subtracting from any value the global mean of its batch; see, for example the *batchadjust* command in the R package PAMR (Hastie et al., 2015). The first linear method that we consider here is calibration by matching each marker’s mean and variance in the source sample to the corresponding values in target sample.

The second common practice is to obtain the principal components of the data, and remove the components that are most correlated with the batch index (Liu and Markatou, 2016). Table 2 compares the performances of our approach and the two approaches mentioned above in terms of MMD scores. As can be seen, the calibration obtained from our MMD-ResNet outperforms the ones obtained by other two methods.

4.3 Calibration of Single-Cell RNA-seq Data

Drop-seq (Macosko et al., 2015) is a novel technique for simultaneous measurement of single-cell mRNA expression levels of all genes of numerous individual cells. Unlike traditional single cell sequencing methods, which can only sequence up to hundreds or a few thousands of cells (Picelli

Table 2: CyTOF calibration: Comparison of calibration using (1) matching mean and variance of each marker, (2) PCA and (3) MMD-ResNet. The table entries are average MMD between the target sample and the calibrated source sample, based on five random subsets of size 1000.

MMD to target \ pair	pa.1 base.	pa.2 base.	pa.1 treat.	pa.2 treat.
mean, var. matching	0.26 \pm 0.02	0.25 \pm 0.01	0.30 \pm 0.01	0.30 \pm 0.02
PCA	0.38 \pm 0.02	0.39 \pm 0.01	0.44 \pm 0.01	0.37 \pm 0.01
ResNet	0.27 \pm 0.01	0.18 \pm 0.01	0.24 \pm 0.01	0.17 \pm 0.01

et al., 2013), (Jaitin et al., 2014), Drop-seq enables researchers to analyze many thousands of cells in parallel, thus offers a better understanding of complex cell types of cellular states.

However, even with several thousands of cells (~ 5000) in each run, only less than half of the cells typically contain enough transcribed genes, that can be used for statistical analysis. As the number of cells in a single run is not sufficient for studying very complicated tissues, one needs to perform multiple runs, in several batches, so that the cumulative number of cells is a good representation of the distribution of cell populations. This process may create batch effects, which need to be removed.

In (Shekhar et al., 2016), seven replicates from two batches were sequenced to study bipolar cells of mouse retina. Applying their approach to clean and filter the data, we obtained a dataset of 13,166 genes, each expressed in more than 30 cells and has a total transcript count of more than 60, and 27,499 cells, each of which has more than 500 expressed genes. Data was then normalized such that counts in each cell sum to 10000, followed by a log transform of (count + 1). Shekhar et al. (2016) estimated that most of the signal is captured by the leading 37 principal components and used them for downstream analysis. We therefore projected the 13,166-dimensional data onto the subspace of the first 37 principal components and used this reduced data for our calibration experiment.

We arbitrarily chose batch 1 to be the target and the one from batch 2 to be the source, and used them to train a MMD-ResNet. The net had three blocks, where each block is as in Figure 2. In each block, the two weight matrices were of size 37×50 and 50×37 . The net weights were initialized by sampling from a $\mathcal{N}(0, 10^{-4})$ distribution. t -SNE plots of the data before and after calibration are presented in Figure 5, which shows that after calibration, clusters from the source batches are mapped onto their target batch counterparts.

Table 3 shows the MMD between the source and target batch before and after calibration, in comparison to the two linear calibration methods mentioned in Section 4.2.5, as well as to Combat (Johnson et al., 2007), a standard technique for batch effect removal, which performs linear adjustments, where the corrections are based on Bayesian estimation. Combat and the mean-variance matching were applied on the full set of 13,166 genes, after normalization as in (Shekhar et al., 2016), rather than on the projection of the data onto the leading 37 principal components, which was the input to the MMD-ResNet on this dataset. As can be seen, MMD-ResNet outperforms all other methods in terms of MMD.

To further assess the quality of calibration, and verify that our approach does not distort

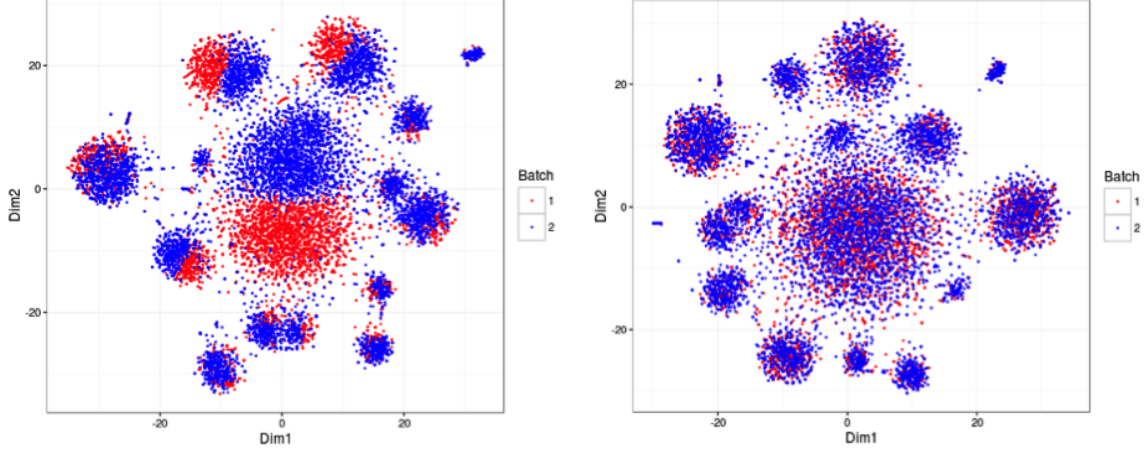


Figure 5: Calibration of scRNA-seq. t -SNE plots before (left) and after (right) calibration using MMD-ResNet.

Table 3: RNA calibration. Comparison of calibration using (1) matching mean and variance of each gene, (2) PCA, (3) Combat and (4) MMD-ResNet. The table entries are average MMD between the target sample and the calibrated source sample, based of five random subsets of size 1000. The MMD between two random batches of the target sample is provided as reference in the rightmost column.

before calib.	mean, var. matching	PCA	Combat	ResNet	target-target
0.43 ± 0.01	0.25 ± 0.01	0.21 ± 0.01	0.15 ± 0.01	0.12 ± 0.01	0.11

the underlying biological patterns in the data, we examine the sub-population of cells with high log-transformed expression values (≥ 3) of the *Prkca* marker (which characterizes the cell sub-population of the large cluster in Figure 5). Figure 6 shows this sub-population before and after calibration, as well as after calibration using Combat. As can be seen, this sub-population is calibrated appropriately. Visually, in this analysis, MMD-ResNet achieves better calibration than Combat.

4.4 Indirect Calibration

In this section we demonstrate how MMD-ResNets can be used to calibrate a source distribution to a target distribution in an indirect manner, i.e., without training a net to learn this map directly, as in the previous experiments. For this experiment we use four of the CyTOF samples described in Section 4.2.1. i.e., samples from patients 1 and 2 at baseline condition, each measured on the instrument in day 1 and day 2. We use the shorthand notation p_1d_1 to refer to the sample of patient 1 measured in day 1 and similarly p_1d_2, p_2d_1, p_2d_2 to the other samples. In Section 4.2.3 we trained a MMD-ResNet (which we now denote by N_{p_1}) that maps




Figure 6: Calibration of cells with high expression of Prkca. *t*-SNE plots before calibration (left), after calibration using Combat (middle) and MMD-ResNet (right).




Figure 7: Indirect calibration experiment scheme.

p_1d_1 to p_1d_2 and a ResNet N_{p_2} which maps p_2d_1 to p_2d_2 . In the following experiment we will map p_1d_1 to p_1d_2 indirectly. The setup is as follows: In addition to the nets N_{p_1}, N_{p_2} that were trained in Section 4.2.3, we train two additional MMD-ResNets, a ResNet N_{d_1} , mapping p_1d_1 to p_2d_1 and a ResNet N_{d_2} , mapping p_2d_2 to p_1d_2 . A scheme showing direct and indirect calibrations is shown in Figure 7.

We then mapped p_1d_1 to p_1d_2 through N_{d_1} , followed by N_{p_2} and N_{d_2} (while adjusting the means and variances at each point, to account for the fact that each of these nets was trained on a standardized source sample), and compared the resulting calibration to the direct calibration obtained by applying N_{p_1} on p_1d_1 . The results are presented in Figure 8. As we can see, the indirect calibration is only slightly less accurate than direct calibration, and removes much of the batch effect. MMD between the source and target values support this observation: while before calibration the MMD is 0.69, it is 0.27 after direct calibration and 0.29 after indirect calibration.

The success of removing much of the batch effect via indirect calibration in the above experiment implies that the biological state of the patient p_1 were not distorted by much during the propagation through the nets N_{d_1} and N_{d_2} . This suggests that our MMD ResNets approach can be adapted for performing calibration in cases where replicates from a reference

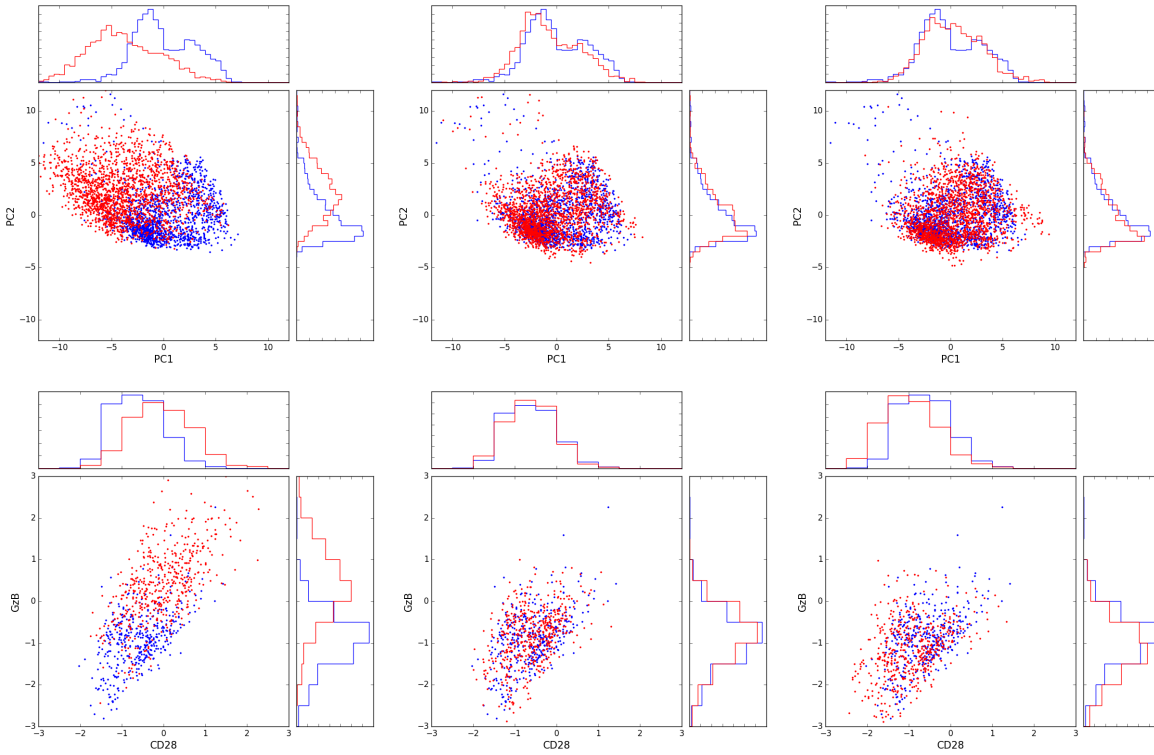


Figure 8: Indirect calibration of CyTOF data. Left: before calibration. Center: direct calibration. Right: indirect calibration. Top row: whole sample, projected onto the subspace of the first two principal components. Bottom row: CD8 sub-population in the (CD28,GzB) plane.

sample are measured in two batches and additional samples are measured only in one of the batches.

Suppose, for example, that in each day we run a CyTOF instrument to measure blood of a several (different) individuals, and in addition we also measure in each of these runs a replicate of a reference blood sample. One can train a MMD-ResNet at each day t to calibrate the reference blood sample to its distribution at day 0. In addition, every replicate measured on day t can be calibrated using (a different) MMD-ResNet to the reference sample at day t . This way data from all days may be compared, by mapping all samples to coordinates of the reference sample at day 0.

5 Related Work

Leek et al. (2010) thoroughly discuss the importance of tackling batch effects and review several existing approaches for doing so.

Bead normalization (Finck et al., 2013) is a specific normalization procedure for CyTOF. As we observed in Section 4, two CyTOF samples may significantly differ in distribution even after Bead normalization. Warping (Hahne et al., 2010) is an approach for calibration of cytometry data where for each marker, the peaks of the marginal distribution in the source sample are (possibly non-linearly) shifted to match the peaks of the corresponding marginal in the target sample. We argue that warping can perhaps be performed by training MMD-ResNet for each single marker. The advantage of MMD-ResNet over a warping is that the former is multivariate, and can take into account dependencies, while the latter assumes that the joint distributions is a product of its marginals (Finak et al., 2014).

Surrogate variable Analysis (Leek and Storey, 2007) is a popular approach for batch effect adjustment, primarily in gene expression data. However, it is designed for supervised scenarios where labels representing the phenotype of each gene expression profile are provided, hence it is not directly applicable

MMD was used as a loss criterion for artificial neural networks in (Li et al., 2015; Dziugaite et al., 2015), where the goal was to learn a generative model that can transform standard input distributions (e.g. white noise) to a target distribution. To the best of our knowledge, MMD nets have not been applied to the problem of removal of batch effects, which is considered here.

6 Discussion

The problem of learning generative models has drawn much attention in the machine learning community recently. Evaluation of such models, however, is not always fully clear. Many recent works proposing generative models use Parzen window estimates for model evaluation. As Theis et al. (2015) nicely point out, evaluation of generative models using Parzen windows is problematic; in our context, for example, suppose that the net maps the source points to the centers of mass of the target sample. Such a map will have high Parzen likelihood estimates, while clearly not calibrating the data well. MMD, which takes also into account the internal

structure of the calibrated source sample (term which is missing in Parzen estimates) might be more suitable for evaluation of the quality of the calibration.

In some of our experiments, which are not reported here, we found out that identifying cluster structure of the data might be a useful practice prior to applying MMD-ResNets in certain applications. For instance, when one uses CyTOF to characterize Peripheral Blood Mononuclear Cells (PBMCs), the multi-marker cell distributions typically have separable clusters, corresponding to cell type sub-populations. While the relative proportion of different cell types in two replicate blood samples is expected to be invariant to the CyTOF machine, measuring these samples in two different runs in the same instrument or two different instruments often show noticeable differences between the cell type composition. When the proportions of corresponding clusters differ between the source and target distributions, we do not expect that MMD-ResNet will account for that difference, as it computes a continuous map. In such cases, for example, it might be useful to use sub-sampling in order to match the relative proportions of each cell type between the source and the target samples.

7 Conclusion

We presented a novel deep learning approach for non-linear removal of batch effects, based on residual networks, to match the distributions of the source and target samples. We applied our approach to CyTOF and scRNA-seq and demonstrated impressive performance. To the best of our knowledge, such a performance on CyTOF data was never reported. Yet, our approach is general and can be applied to various data types. To justify our approach, we showed that equivalent nets that lack the shortcut connections may distort the biological conditions manifested in the samples, while residual nets preserve them. We also presented a novel approach for indirect calibration, which, to the best of our knowledge, is not performed elsewhere. It is based on an appealing property of using neural nets for calibration, which is the fact that the nets define a map, that can be later one applied to new data.

Lastly, despite the impressive experimental results presented here, a two sample test (say, a permutation test using MMD as a test statistic) will reject the hypothesis that the calibrated source sample has the same distribution as the target sample. Yet, in the same way that general deep learning techniques, operating on raw data outperform traditional algorithms tailored for specific data types and involving domain knowledge and massive pre-processing, we find our proposed approach and experimental results very promising and hope that they open new directions for removing batch effects in biological datasets. For example, recent proposed experimental approaches to standardization (Kleinstuber et al., 2016), should provide an excellent source for application of MMD-ResNet for calibration.

Acknowledgement

This research was partially funded by NIH grant 1R01HG008383-01A1 (Y.K.).

References

- Dziugaite, G. K., Roy, D. M., and Ghahramani, Z. (2015). Training generative neural networks via maximum mean discrepancy optimization. *arXiv preprint arXiv:1505.03906*.
- Finak, G., Jiang, W., Krouse, K., Wei, C., Sanz, I., Phippard, D., Asare, A., Rosa, S. C., Self, S., and Gottardo, R. (2014). High-throughput flow cytometry data normalization for clinical trials. *Cytometry Part A*, 85(3):277–286.
- Finck, R., Simonds, E. F., Jager, A., Krishnaswamy, S., Sachs, K., Fantl, W., Pe’er, D., Nolan, G. P., and Bendall, S. C. (2013). Normalization of mass cytometry data with bead standards. *Cytometry Part A*, 83(5):483–494.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Aistats*, volume 9, pages 249–256.
- Gretton, A., Borgwardt, K. M., Rasch, M., Schölkopf, B., and Smola, A. J. (2006). A kernel method for the two-sample-problem. In *Advances in neural information processing systems*, pages 513–520.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. (2012). A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773.
- Hahne, F., Khodabakhshi, A. H., Bashashati, A., Wong, C.-J., Gascoyne, R. D., Weng, A. P., Seyfert-Margolis, V., Bourcier, K., Asare, A., Lumley, T., et al. (2010). Per-channel basis normalization methods for flow cytometry data. *Cytometry Part A*, 77(2):121–131.
- Hardt, M. and Ma, T. (2016). Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*.
- Hastie, T., Tibshirani, R., Narasimhan, B., Chu, G., and Tibshirani, M. R. (2015). Package pamr.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Identity mappings in deep residual networks. *arXiv preprint arXiv:1603.05027*.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- Jaitin, D. A., Kenigsberg, E., Keren-Shaul, H., Elefant, N., Paul, F., Zaretsky, I., Mildner, A., Cohen, N., Jung, S., Tanay, A., et al. (2014). Massively parallel single-cell rna-seq for marker-free decomposition of tissues into cell types. *Science*, 343(6172):776–779.
- Johnson, W. E., Li, C., and Rabinovic, A. (2007). Adjusting batch effects in microarray expression data using empirical bayes methods. *Biostatistics*, 8(1):118–127.

- Kleinstaubner, K., Corleis, B., Rashidi, N., Nchinda, N., Lisanti, A., Cho, J. L., Medoff, B. D., Kwon, D., and Walker, B. D. (2016). Standardization and quality control for high-dimensional mass cytometry studies of human samples. *Cytometry Part A*.
- Leek, J. T., Scharpf, R. B., Bravo, H. C., Simcha, D., Langmead, B., Johnson, W. E., Geman, D., Baggerly, K., and Irizarry, R. A. (2010). Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics*, 11(10):733–739.
- Leek, J. T. and Storey, J. D. (2007). Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genet*, 3(9):e161.
- Li, S., Jiao, J., Han, Y., and Weissman, T. (2016). Demystifying resnet. *arXiv preprint arXiv:1611.01186*.
- Li, Y., Swersky, K., and Zemel, R. (2015). Generative moment matching networks. In *International Conference on Machine Learning*, pages 1718–1727.
- Liu, Q. and Markatou, M. (2016). Evaluation of methods in removing batch effects on rna-seq data. *Infectious Diseases and Translational Medicine*, 2(1):3–9.
- Macosko, E. Z., Basu, A., Satija, R., Nemesh, J., Shekhar, K., Goldman, M., Tirosh, I., Bialas, A. R., Kamitaki, N., Martersteck, E. M., et al. (2015). Highly parallel genome-wide expression profiling of individual cells using nanoliter droplets. *Cell*, 161(5):1202–1214.
- Nygaard, V., Rødland, E. A., and Hovig, E. (2016). Methods that remove batch effects while retaining group differences may lead to exaggerated confidence in downstream analyses. *Biostatistics*, 17(1):29–39.
- Picelli, S., Björklund, Å. K., Faridani, O. R., Sagasser, S., Winberg, G., and Sandberg, R. (2013). Smart-seq2 for sensitive full-length transcriptome profiling in single cells. *Nature methods*, 10(11):1096–1098.
- Shekhar, K., Lapan, S. W., Whitney, I. E., Tran, N. M., Macosko, E. Z., Kowalczyk, M., Adiconis, X., Levin, J. Z., Nemesh, J., Goldman, M., et al. (2016). Comprehensive classification of retinal bipolar neurons by single-cell transcriptomics. *Cell*, 166(5):1308–1323.
- Spitzer, M. H. and Nolan, G. P. (2016). Mass cytometry: Single cells, many features. *Cell*, 165(4):780–791.
- Theis, L., Oord, A. v. d., and Bethge, M. (2015). A note on the evaluation of generative models. *arXiv preprint arXiv:1511.01844*.
- Tieleman, T. and Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4(2).
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. (2008). Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM.

A Specification of Markers in CyTOF Experiments

Table 4 provides the information about the 25 markers used in the CyTOF experiments in Section 4.2.

B Additional Plots for CyTOF calibration

Figure 9 shows the projection of the source and target samples onto the first two principal components of the target sample for the three additional source-target pairs not shown in Figure 1.

Figure 10 shows the projection of the CD8+ T-cell sub-population in the source and target data onto the first two principal components of the target sample for the three additional source-target pairs not shown in Figure 4.

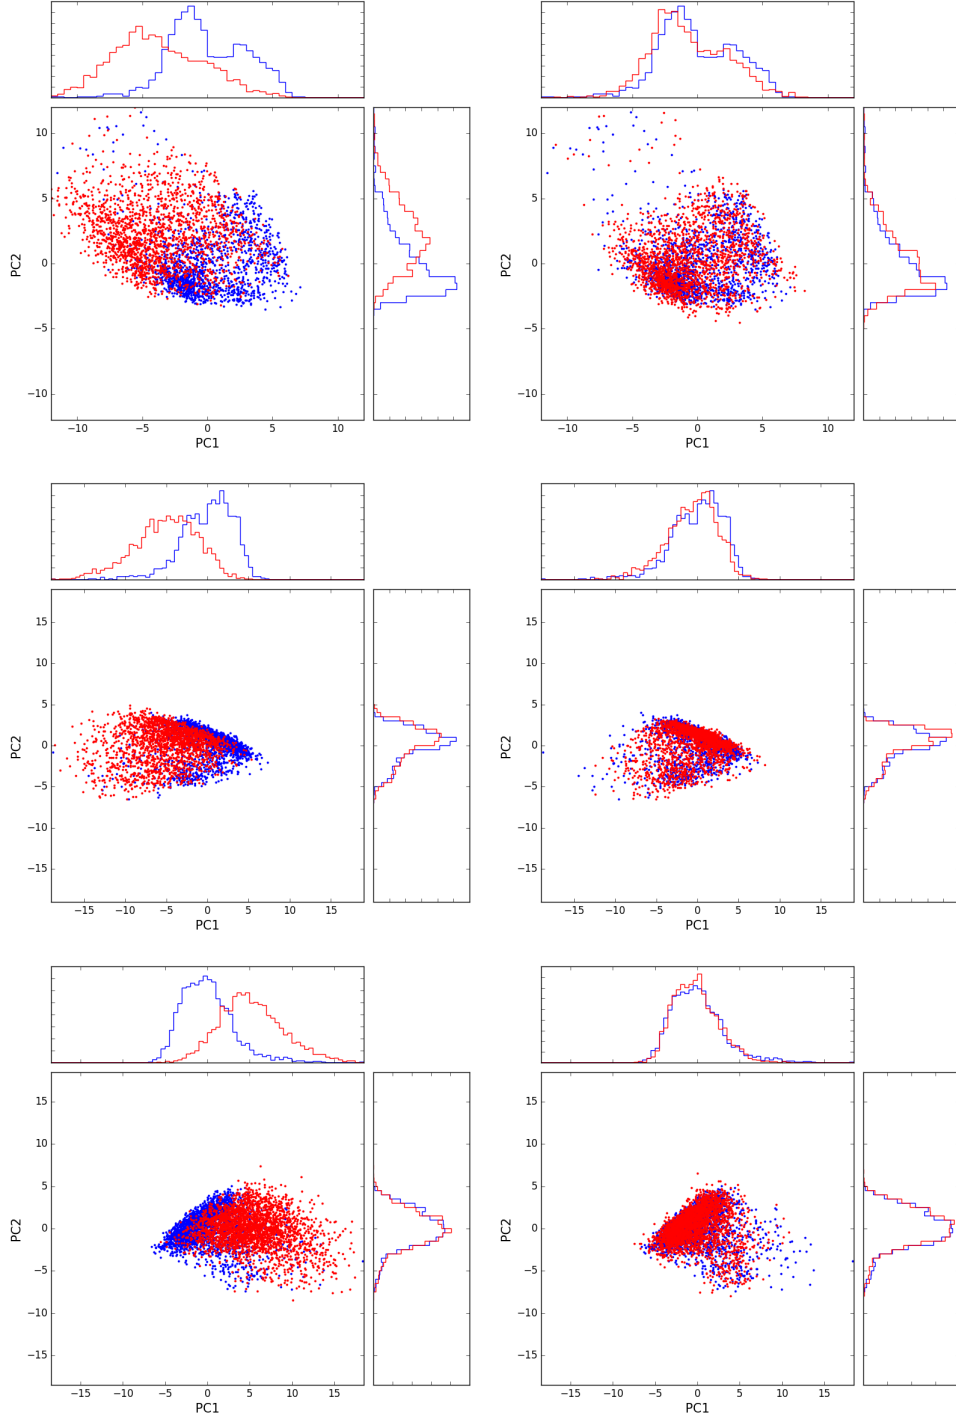


Figure 9: Calibration of CyTOF data, for each of the three source-target pairs not shown in Figure 1. Projection of the source (red) and target (blue) samples on the first two principal components of the target data. Left: before calibration. Right: after calibration.

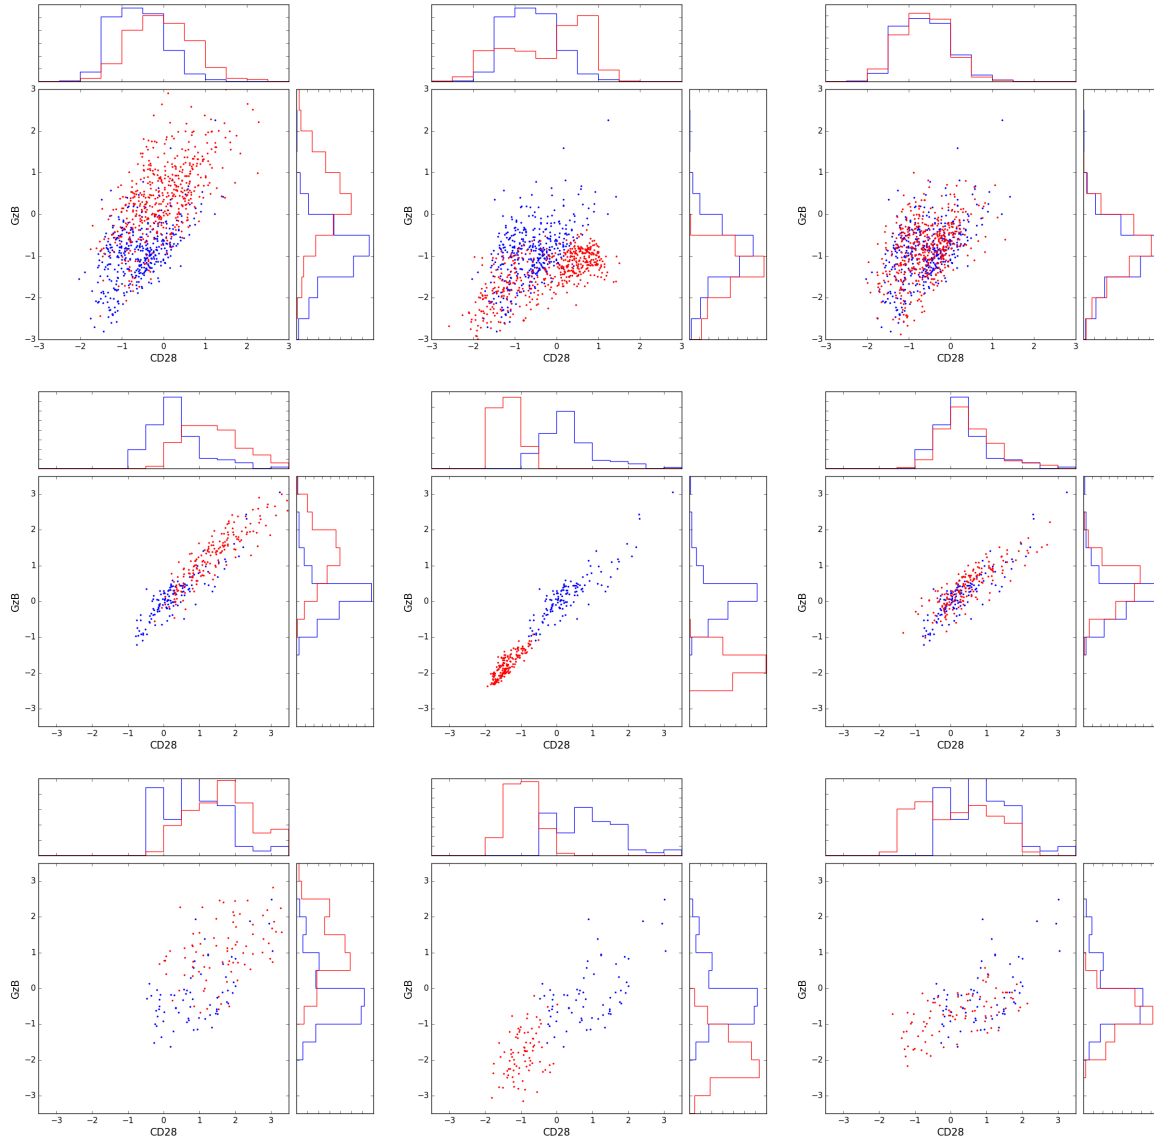


Figure 10: Calibration of CD8+T-cells sub-population in the (CD28,GzB) plane, for each of the three source-target pairs not shown in Figure 4. In each row the left plot corresponds to before calibration, the right to calibration using ResNet, and the center to calibration using an identical net, without shortcut connections and initialized in a standard fashion.

Table 4: Specification of the 25 markers used to characterize cell sub-populations in our CyTOF experiments. Bead standards are embedded in each sample to allow Bead normalization. Each Bead contains the four heavy metal isotopes labeled by 1 in the third column.

Isotope	Marker	Beads
89Y	CD45	0
142Nd	CD19	0
143Nd	CD127	0
145Nd	CD4	0
146Nd	CD8a	0
147Sm	CD20	0
149Sm	CD25	0
151Eu	CD278	1
152Sm	TNFA	0
153Eu	Tim3	1
155Gd	CD27	0
156Gd	CD14	0
159Tb	CCR7	0
160Gd	CD28	0
161Dy	CD152	0
162Dy	FOXP3	0
164Dy	CD45RO	0
165Ho	INFg	1
166Er	CD223	0
167Er	GzB	0
170Er	CD3	0
172Yb	CD274	0
174Yb	HLADR	0
175Lu	PD1	1
209Bi	CD11b	0